



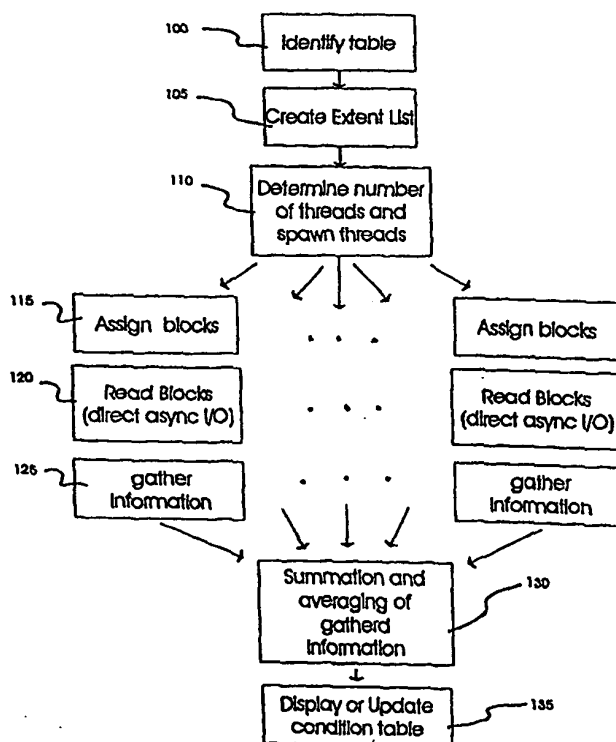
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/30		(11) International Publication Number: WO 00/13111
A1		(43) International Publication Date: 9 March 2000 (09.03.00)
(21) International Application Number: PCT/US99/19853 (22) International Filing Date: 30 August 1999 (30.08.99) (30) Priority Data: 09/144,011 31 August 1998 (31.08.98) US (71) Applicant: PLATINUM TECHNOLOGY IP, INC. [US/US]; 1815 South Meyers Road, Oakbrook Terrace, IL 60181 (US). (72) Inventor: PEREIRA, Hilton, M.; 3092 Saffron Circle, Thousand Oaks, CA 91360 (US). (74) Agents: FLIESLER, Martin, C. et al.; Fliesler Dubb Meyer and Lovejoy LLP, Suite 400, Four Embarcadero Center, San Francisco, CA 94111-4156 (US).		(81) Designated States: JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published With international search report.

(54) Title: METHOD AND APPARATUS FOR FAST AND COMPREHENSIVE DBMS ANALYSIS

(57) Abstract

A fast and comprehensive analysis of a database table is performed by reading a header block describing the location of blocks storing data information of a database table. The data blocks are read using direct asynchronous I/O into memory (120). Data read in from the data blocks is analyzed to determine information regarding the health or condition of the database table. The analysis is performed by spawning separate processes (110), each process being assigned an approximately equivalent number of data blocks to analyze (115). Once each process completes analysis of the assigned data blocks, the analysis by each thread is summarized and averaged (130) to provide information to the DBA relating to the condition of the database table (135).



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

- 1 -

**Method and Apparatus for Fast and Comprehensive
DBMS Analysis**

Background of the Invention

5 Field of the Invention

 This invention relates to a method of analysis and of increasing the speed at which the condition of a database is determined. The invention is more particularly related to the determination of appropriate data needed for a comprehensive analysis of specific tables in the database, and to increase the speed at which the analysis for determining the condition of a database table is performed.

15 Discussion of the Background

 With the increasing use of computer databases in each of the home, office, and scientific environments, computer users and Database Administrators (DBA's) are increasingly confronted with the tasks of determining the condition or health of databases and performing various operations to maintain the database, including reorganization or rebuilding of tables stored in the database. Before rebuilding a database table(s), the DBA must first determine whether the condition of the database table(s) requires rebuilding.

 Typically, commercial database products include tools to perform analysis to determine the condition of a database. For example, Oracle database management systems include an analyze function that performs a predetermined function including an SQL call UPDATE_STATISTICS_FOR_TABLE. The analyze function reads a selected database table and generates a set of statistics that reveal the condition of the table.

- 2 -

Figure 1 is a flowchart illustrating the basic steps performed by a conventional Database Management System (DBMS) for determining the condition of a database table. At step 10, a Database Administrator (DBA) provides the name of a table to be analyzed. The table name is provided to a DBMS function that returns data about the named table (step 20). The returned data is analyzed resulting in statistical information (step 30) which is then placed (stored) in a location for retrieval for by DBA (step 40).

The DBMS function that returns data (DBMS analysis function) is shown in Fig. 2 and includes the steps of building an SQL call (step 50), invoking the SQL call (step 60), performing SQL overhead and administrative functions (step 70), and returning the SQL results (step 80). The SQL results are provided as the return values (statistical information) in step 30 of Fig. 1.

Once the SQL results are analyzed, the displayed (or stored) information is utilized by the DBA to determine appropriate action, if needed, to properly maintain the database (rebuild tables, for example). In the example function, UPDATE_STATISTICS_FOR_TABLE, the Oracle DBMS provides table information including No. of Rows, Chained Rows, Blocks allocated, Blocks allocated but free, Blocks on freelist, Average available space in Freelist, Average free space in a Block, and Min, Max, Average Row Length.

However, the provided table information does not provide enough information to give the DBA the most accurate picture of the DBMS table condition. In addition, current methods for retrieving the DBMS table information are slow, being encumbered by overhead

- 3 -

associated with standard DBMS procedures and function calls (file reading techniques, and SQL, for example).

Modern database tables contain enormous amounts of data. If a DBA needs to determine the condition of a large database table, analysis techniques based on current methods can take unacceptably large amounts of time during which a database table may not be available to a user. Furthermore, current methods do not provide the DBA with the exact information needed for a proper analysis. Therefore, the DBA may rebuild tables when it is not necessary or where other action may have been more appropriate. All of the above leads to increased down time and maintenance efforts.

Summary of the Invention

The present inventor has realized the need for increasing the amount of data gathered to determine the condition of tables within a database and to increase the speed at which the condition is determined. Accordingly, it is an object of the present invention to provide a comprehensive analysis of at least one database table to determine an information set relating to the health of the tables analyzed.

It is another object of this invention to provide method to increase the speed at which database tables are analyzed.

These and other objects are accomplished by a method for analyzing database tables including the steps of identifying a database table, determining elements corresponding to the condition of the database table based on data contained in the database tables and providing said elements to one of a display for viewing by a user, a storage location, and a mechanism

- 4 -

for using said elements. The elements determined include at least one of number of simple rows, number of migrated rows, number of deleted rows, number of blocks initialized but empty, average free space in a block unused, average free released free space, average row length of migrated rows, average row length of chained rows, number of blocks to be read for normal, chained, and migrated rows, measure of density of blocks used, number of blocks exceeding a percentage used threshold, number of blocks less a percentage free threshold, number of chained rows spanning data files, and min max average of chained rows spanning blocks.

In addition, a method to increase the speed of analysis is provided including the steps of identifying a table in said database, reading each of blocks associated with said table using direct asynchronous I/O, and performing said analysis based on information contained in said blocks.

Additional speed of analysis is provided by parallel processing and bypassing SQL overhead associated with regular RDBMS architectures. The parallel processing includes the steps of dividing the blocks read into plural sets, and providing each set to separate processes to determine information needed for the analysis.

Brief Description of the Drawings

A more complete appreciation of the invention and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the following detailed description when considered in connection with the accompanying drawings, wherein:

- 5 -

Figure 1 is a flow chart of a conventional routine for determining status of a database table;

Figure 2 is a flow chart of a conventional database management system (DBMS) analysis function for retrieving information needed to determine the status of a database table;

Figure 3 is an illustration of processing of a data block; and

Figure 4 is a flowchart of the Fast Analysis method for accessing a database table and analyzing table contents according to the present invention.

Detailed Description of the Invention

Conventional Database Management Systems (DBMS) utilize tables to store related data in a database. Once a table is defined, the DBMS will allocate blocks for storage of the related data. For example, if a table is to be built to contain 100,000 rows, and the rows of the table were defined such that 100 rows could fit per block, the DBMS would allocate 1,000 blocks to build the table.

Generally, DBMS systems allocate blocks in sets of contiguous blocks. A contiguous set of allocated blocks is commonly referred to as an extent. As a general rule, extents vary in size. Using the above example, the DBMS may utilize a single extent of 1,000 blocks, 2 extents of 500 blocks, or any other combination of extent sites to allocate the required 1,000 blocks to build the table. Once the required blocks are allocated, the data is then stored in rows in the table utilizing the blocks allocated.

Over the course of time, additional data may be added to the table and the DBMS will allocate

- 6 -

additional blocks as required. For example, if a user adds 250 rows to the table, using the above parameters, an additional 3 blocks would need to be allocated. Also over the course of time, information in the database may be deleted. In this case, deletion of rows occurs through the use of SQL to delete rows from the table. For example, a user may delete fifty rows from block 1, 40 rows from block 20, and 30 rows from block 60. When this occurs, it causes there to be more blocks in the table than required to maintain the data present in the table.

From time to time, the Database Administrator (DBA) will perform an analysis on each table that provides information regarding the condition of the database table. For example, the database administrator may look at information regarding the number of deleted rows to ascertain efficiency information with respect to how many blocks are holding deleted rows.

As another example, the Database Administrator may look to see whether a table has become fragmented. Row migration is one type of fragmentation. Row migration occurs when a row gets updated and the original space allocated for the row in the block is insufficient to hold all of the updated information. In this case, the row is moved to a location with more space, and a pointer is placed in the block where the original row resided, the pointer now pointing to the location of the moved (migrated) row. A large amount of row migration would tell the DBA that perhaps an inadequate amount of space was allocated for the original rows.

More often than not, there is insufficient space in a block to hold an updated row. In this case, the

- 7 -

row is migrated to an entirely different block than where the original row resided, and the pointer is placed in the original row position points to the updated row in the different block.

5 Row migration causes increased retrieval time of database information because a pointer must be read and interpreted, and possibly two or more blocks (the block containing the pointer, and the block containing the actual row data) must be read to retrieve row specific
10 information, as compared to a single block read for tables having no row migration.

 It is useful for the database administrator to be able to retrieve information on row migration within a database table. If a lot of row migration has occurred,
15 the database administrator may likely decide to rebuild the tables to more efficiently allocate row space.

 In another example, a DBA makes a decision regarding the management of a row allocation within blocks of a table by setting a percentage of blocks
20 free (PCTFREE) or percentage of blocks used (PCTUSED). For example, the DBMS allows the DBA to set PCTFREE when a database table is created. Then, as the DBMS fills up each block with table information, it will keep a percentage of the block free at least equal to
25 the percentage PCTFREE for row updates.

 The DBA sets the PCTFREE variable depending on how the database table is to be used. For example, if a table is to have frequent updates, additional PCTFREE would be established so that enough space is available
30 to allow any necessary row migration to occur within the same block. As discussed above, row migration within the same block causes a table to become fragmented. However, rows migrated within the same

- 8 -

block are retrieved with a single block read rather than the cumbersome process of reading a block, interpreting a pointer, and reading a second block (or more) as occurs when rows are migrated to different blocks. Therefore, appropriate PCTFREE settings allow DBMS performance to be maintained although the database tables may be modified.

PCTUSED is another parameter that allows the DBA to control the row allocation of a DBMS table. The DBMS prevents additional rows to be placed in a block unless the percentage of that block has fallen below PCTUSED. PCTUSED is different from PCTFREE in the sense that although a block may be used if there is an update, it will not be used to insert a new row unless the percentage used in the block is below PCTUSED. Therefore, PCTFREE is used to maintain free space when the table is initially built, and PCTUSED determines whether a new row may be inserted in a block. The PCTUSED parameter comes into effect when used space becomes free due to deletes.

In analyzing the database table, a DBA will look to see how closely the use of blocks of a table correspond to PCTUSED and PCTFREE set for that table. If the average free space for a table is high, the DBA may decide to rebuild the table with a lower PCTFREE because space is being wasted. Conversely, if the hereinbelow described parameter BOCKS_LESS_THAN_PCTUSED is sufficiently high, it tells the DBA that too much migration is taking place and that it may be preferable to rebuild the tables with optimal values for PCTFREE and/or PCTUSED.

Thus, the value of knowledge regarding the condition of database tables is firmly established.

- 9 -

However, conventional database management systems do not provide enough information for a DBA to make the most intelligent and efficient decisions regarding how and when tables within a database should be altered, modified, or rebuilt.

The present invention, hereinafter referred to as Fast Analysis, provides a detailed view of the condition, or health, of the database table allowing the DBA to make well informed decisions as to how to maintain the DBMS tables.

Referring to Table 1, a chart is provided comparing information gathered by Fast Analysis in comparison to a conventional DBMS system (Oracle, in this example). As shown in Table 1, Fast Analysis provides information regarding the condition of database tables in greater depth and quantity as compared to the conventional system.

Table 1: Comparison of Fast Analysis and Oracle Information Gathered

Name	Fast Analy sis	Ora cle	Description
ROW COUNT	X	X	No. of Rows
NORMAL ROW COUNT	X		No. of Simple Rows
MIGRATE ROW COUNT	X		No. of Migrated Rows
CHAINED ROWS	X	X	No. of Chained Rows
DELETED ROWS	X		No. of Deleted Rows

- 10 -

5	BLOCK COUNT	X	X	No. of Blocks allocated for table
	USED BLOCK COUNT	X	X	No. of Blocks used by a table
	EMPTY BLOCK COUNT	X		No. of Blocks initialized but empty
	FREE BLOCK COUNT	X	X	No. of Blocks allocated but free
	BLOCKS ON FREELIST	X	X	No. of Blocks on the freelist
10	AVG SPACE IN FREELIST BLKS	X	X	Average available space in freelist blocks
	AVG FREE SPACE	X	X	Average free space [across blocks in a table] in a block
	AVG UNUSED FREE SPACE	X		Average free space in a block never used
15	AVG RELEASED FREE SPC	X		Average free released free space in box
	ROW LENGTH	X	X	Min,Max,Average Row length
	AVG MIGRATED ROW LEN	X		Average row length of chained rows
20	BLOCK READS	X		No. of Blocks to be read for normal, chained and migrated rows.
	BLOCK FILL PERCENTAGE	X		Measure of Density of the blocks used
	BLKS MORE THAN PCTFREE	X		No. of Blks exceeded PCTFREE
25	BLOCKS LESS THAN PCTUSED	X		No. of Blks less than PCTUSED
	FILESPAN COUNT	X		No. of chained rows spanning datafiles
	BLOCKSPAN OFFSET	X		min,max average of chained rows spanning blocks

- 11 -

	AVG BLOCK HEADER	X		Avg. Block header size
	AVG BLOCK DATA	X		Avg. space available in a block
5	MIGRATED AND CHAINED CNT	X		No. of rows migrated and chained
	MIGRATED OR CHAINED CNT	X		No. of rows either migrated or chained
	AVG ROW DIR ENTRIES	X		Reusable offsets to row data
10	TOTAL ROW DIR ENTRIES	X		No. of row directory entries
	AVG NORMAL ROW LEN	X		Avg. length of a normal row
15	AVG CHAINED ROW LEN	X		Avg. length of migrated and chained rows
	AVG MIGRATED AND CHAINED ROW LEN	X		Avg. length of migrated and chained rows
	HOME ROW READS	X		No of initial blocks to read all rows of a table
20	MIGRATED ROW READS	X		No of additional blocks to read for migrated rows
	CHAINED ROW READS	X		No of additional blocks to read for chained rows

25

In order to determine the information in Table 1,
 the structure of the database tables is first
 established. For example, in an Oracle database, table
 structure is based on an Oracle block which is defined
 as a unit of data storage (e.g., an Oracle block
 storing rows of data for a table is an Oracle data
 block). Block size in bytes is defined by an instance
 parameter by the DBA at the time of configuration and

30

- 12 -

different types of Oracle blocks store different types of information.

As the tables are built, groups of blocks, known as extents, are allocated. Every Oracle table has an
5 initial allocation of blocks which constitutes the first extent for an Oracle table. Subsequent extents (contiguous sets of blocks) are allocated to accommodate growth of the table and are allocated based on a next extent parameter specified at table creation.

10 Every table has an extent segment header which maintains an extent list in the form of file number (identifying a file), block number (identifying an offset within the file), and length for each extent of the table. In Oracle, the extent segment header is the
15 starting block for a given table. The extent list and a free list is stored in this block. The address of this block (extent segment header) is available as part of information stored in a table which is part of a system dictionary (Oracle system dictionary, for
20 example).

In Oracle, a unique block address is a combination of a file number and a block number. The file number is a unique number that is assigned to each data file managed by the DBMS. A block number is relative to a
25 file and relates to an offset in the file. Therefore, a file offset can be calculated given the block number and the block size for the instance.

Although the above structural elements (blocks, extents, extent segment header, etc.) are described in
30 terms of an Oracle database and tables, the same structure applies to other databases and tables as well (Sybase and Informix, for example). In addition, other

- 13 -

database structures can also be appropriately described utilizing the same or related terminology.

For example, although a block is a typical unit data storage, other database implementations may be based on pages. In this case, an extent would be a group of pages rather than blocks, and the extent itself may be described using other terminology. Therefore, an extent can be generically described as a group of contiguous pages, blocks, or any unit of data storage. In addition, the extent segment header or extent segment header block can be generically defined as any file or other storage where header information about a table can be found, i.e., any location or locations where various extents and their respective sizes used in building one or more database tables may be considered an extent segment header.

Consistent with the above discussion, it is important to note that although Fast Analyzer as described herein provides information described in terms of specific parameters consistent with the structure of an Oracle database (blocks, for example), these terms are applicable in a generic sense to all database applications and therefore Fast Analyzer should not be limited to any specific database implementation.

The following is an element by element explanation of the table information gathered by Fast Analyzer. Each piece of information (element) is defined along with a clarifying explanation including implications where applicable. Information gathered by Fast Analyzer are typically stored in a repository, such as the DAO repository, for example, as defined hereinbelow. The information gathered by Fast Analyzer

- 14 -

is divided into the following categories: block level information, block space usage information, row level information, row level space information, and information to measure block and file I/O.

5

Block Level Information

1. USED_BLOCK_CNT

10 Total number of Blocks that have one or more rows for a given Oracle table. This may be less than the number of blocks allocated for the table.

2. EMPTY_BLOCK_CNT

15 Total number of Blocks that are included in the extent and block high water mark, but have no rows in them. The blocks are initialized for use and are potentially the next candidates to be filled when new rows are inserted for the given table. The extent and block high water indicators are stored in the segment header block of an Oracle table to indicate the current allocation usage of an Oracle table. Sum of
20 USED_BLOCK_CNT and EMPTY_BLOCK_CNT is the value for BLOCKS in the DBA_TABLES view.

25 3. FREE_BLOCK_CNT

 Total number of Blocks which are part of the allocation for the table but have never been initialized. Same as the Oracle's EMPTY_BLOCKS in the DBA_TABLES view.

30

4. EMB_CNT

 Total number of blocks used by the segment header to hold an unlimited extent list. When a table is

- 15 -

created with the unlimited extents option, it is possible that the list of extents which are stored in the segment header exceed the space available in a block, the extent segment header entries points to
5 additional blocks which contain extent list and or freelist information.

5. FREELIST_BLOCK_CNT

Total number of blocks on the freelist. Oracle
10 maintains a list of blocks available for new row inserts. If an Oracle block is filled less than PCT_FREE this block is on the freelist. If rows in an Oracle block get deleted after reaching the PCT_FREE limit and the block usage falls below the PCT_USED
15 limit the block is put on the freelist.

BLOCK SPACE USAGE INFORMATION

1. AVG_RELEASED_FREE_SPACE

20 Released free space in a block is defined as the space previously used and now available due to row deletions or migrations. Released free space need not be contiguous in the block. Average released free space is the average of the released free space in
25 bytes for the blocks of the table.

2. AVG_UNUSED_FREE_SPACE

Unused free space is defined as the free space in the block not yet used by the block for any rows.
30 Unused free space is contiguous, given the fact Oracle stores rows bottom up, it is the space between the row directory and the last row inserted. Average unused

- 16 -

free space is the average of the unused free space in bytes of the blocks for the given table.

3. **AVG_FREE_SPACE**

5 Free space is the sum of the unused free space and the released free space for a block. Average free space in bytes, is the average of the free space available in the blocks for the given table. This statistic is the same as the AVGSPC statistic in Oracle 7.x and AVG_SPACE statistic in Oracle 8.x found in the
10 DBA_TABLES view.

4. **AVG_SPACE_FREELIST_BLOCKS**

15 Average Free space in bytes found in Oracle blocks that are on the freelist. This statistic is the same as the AVG_SPACE_FREELIST_BLOCKS in Oracle 8.x found in the DBA_TABLES view. This information is not reported by Oracle in versions prior to 8.x.

20 5. **USED_BLOCK_FILL_PERCENTAGE**

 This statistic is a measure of the block fill for blocks that have one or more rows in them. The percentage is measure of the amount of space used for data in a block versus the amount of space available
25 for data in a block, not the block size.

6. **BLOCKS_FILL_MORE_PCTFREE**

 The number of blocks which are found to be filled more than the PCT_FREE limit assigned to the table.
30 This occurs when rows in the block are updated and get migrated within the block.

- 17 -

7. BLOCK_FILL_LESS_PCTUSED

The number of blocks which are found to be empty less than the PCT_USED limit assigned to the table. This occurs when rows in the block are either deleted or migrated and the space occupied by rows falls below the PCT_USED limit.

8. AVG_BLOCK_HEADER_BYTES

Each data block contains a block header which contains the table directory, row directory and transaction information storage areas. The AVG_BLOCK_HEADER_BYTES is the average space used in bytes to store block header information for the Oracle blocks of a given table.

9. AVG_BLOCK_DATA_BYTES

The space available for row data storage in a block. The Average space available in bytes for row data storage in the Oracle blocks of a given table.

ROW LEVEL INFORMATION**1. ROW_CNT**

The total number of rows found in the blocks for the given table. This statistic is the same as the NUM_ROWS statistic in Oracle, found in the DBA_TABLES view.

2. NORMAL_ROW_CNT

The total number of rows that are neither migrated nor chained and are found as one contiguous piece in an Oracle block.

- 18 -

3. MIGRATE_ROW_CNT

The total number of rows that have been found to be migrated for a given table. Migration of a row normally occurs when a row is being updated cannot be updated in place, hence Oracle stores a forwarding pointer and relocates the entire row in the same or new block.

4. CHAIN_ROW_CNT

The total number of rows that are chained for a given table. A row is defined to be chained when the entire row does not fit into an Oracle block and hence is stored as row pieces in separate blocks, each row piece has forward and backward pointers linking the row pieces together.

5. MIGRATE_AND_CHAIN_ROW_CNT

The total number of rows that are migrated and chained as per the previous definitions

6. CHAINED_OR_MIGRATED

The sum of all rows that are either chained or migrated or both. Oracle reports this statistic as CHAIN_CNT in the DBA_TABLES view.

7. DELETE_ROW_CNT

The total number of rows that are marked as being deleted in the blocks for a given table. A row when deleted is just marked as being deleted in the row directory of an Oracle block and the space occupied is not reclaimed. The DELETE_ROW_CNT is a count of such entries.

- 19 -

8. AVG_ROW_DIR_ENTRIES

Row directory entries are entries in an Oracle block that contain offsets that point to row data. These entries are created as new rows are inserted but
5 never deleted. These entries may be reused.

9. TOTAL_ROW_DIR_ENTRIES

The Total number of row directory entries found in all the Oracle blocks for a given table. This number
10 may be more than the number of rows if deletions and updates occur on the given table.

ROW LEVEL SPACE INFORMATION

15 1. AVG_ROW_LEN, MIN_ROW_LEN, MAX_ROW_LEN

The average, minimum and maximum row length in bytes found in the Oracle Blocks for a given table. The AVG_ROW_LEN statistic is the same as AVG_ROW_LEN Oracle reported statistic in the DBA_TABLES view for a
20 given table.

2. AVG_NORMAL_ROW_BYTES

The average length in bytes for a normal rows of a table, see normal row definition.
25

3. AVG_MIGRATED_ROW_BYTES

The average length in bytes for a migrated rows of a table, see migrated row definition.

30 4. AVG_CHAINED_ROW_BYTES

The average length in bytes for a chained rows of a table, see chained row definition.

- 20 -

5. AVG_MIG_AND_CHAINED_ROW_BYTES

The average length in bytes for migrated and chained rows of a table, see migrated and chained row definition.

5

INFORMATION TO MEASURE BLOCK AND FILE I/O**1. MIN, AVG, MAX CHAIN_BLOCKSPAN_OFFSET**

Chained rows span blocks, in an Oracle data file block numbers are relative to the start of the data file. The blockspan offset is a measure of the distance between blocks containing chained row pieces. The minimum, average and maximum is collected for a given Oracle table that has chained rows.

15

2. CHAINED_ROWS_THAT_SPAN_FILES

The total number of chained rows that have row pieces in more than one Oracle datafile.

20 3. HOME_ROW_READS

The total number of initial blocks to be read if all the rows of the table were requested for. If the table did not have any migrated or chained rows then the total number of blocks to read to access all rows would be equal to the HOME_ROW_READS.

25

4. MIGRATED_ROW_READS

The total number of additional blocks to be read if all migrated rows of the table were requested for. A row can be migrated more than once and it is possible to have more than two block i/o's for one row.

30

- 21 -

5. CHAINED_ROW_READS

The total number of additional blocks to be read if all chained rows of the table were requested for. A row can be chained over more than one block and it is possible to have more than a two block i/o for one row.

The above described information provides the DBA with a comprehensive set of data for determining the condition of a database table. Each element is determined by reading the blocks of a database table, examining the contents of the blocks, and deriving the information needed to determine the element.

Proper examination of the contents of a block requires knowledge regarding the structure of the block. Block structure may be obtained either from a specification describing the structure or by investigation of a block already created. For example, Fig. 3 illustrates the structure of an Oracle data block. Each data block of an Oracle table is made up of a block header 90, block transaction information area 92, and a table and row directory 94 containing pointers to row data 96. The row data itself is filled bottom up.

The above described elements are divided into block level information, block space usage information, row level information, row level space information, and information to measure block and file I/O, each present in some form or another in the block itself. For example, block space usage information is present in the block header 90. The number of rows, number of row directory entries, and deleted rows for the block can be found in the table and row directory area 94. The

- 22 -

table and row directory provides pointers to find actual row data.

Row types to determine migrated, chained, and normal rows can be decoded from a row type that is stored within each row. Each row is looked at to compute row size which provide the row level space information. Within each row is provided pointers if the row is incomplete (i.e. is migrated or chained). Following these rows is necessary to determine actual row size information and related elements. Block level information is typically gathered from the extent segment header for the table and information provided in the block header itself.

Most of the information are computed in various ways to enable future modeling. Therefore, FASTANAL must collect the following for most metrics:

- * Minimum
- * Maximum
- * Average
- * Standard Deviation (Optional)

Minimum and Maximum are determined for each piece of information (metric). The easiest way to compute the Average is to keep a running total of the sum of the metric, as well as a count of the number of data-points collected. Then, use the formula below for the standard mean (Average):

$$A = X / n$$

where:

X=data value

n=number of points in each series

Standard Deviation is then computed with a second pass by comparing the deviation of each data value with the Average computed above. Since Standard Deviation

- 23 -

requires a second pass, it must be specifically requested.

$$\sigma = \sqrt{\frac{\sum (X-A)^2}{n}}$$

5 The present inventor has also realized that, in addition to extending the analysis performed on database tables, there is a need to improve the speed of the analysis. Fast Analysis provides an extended information set (metrics), and improves the speed at
10 which the metrics are retrieved and processed.

 Conventional database management systems utilize SQL or related commands to provide information for a limited analysis of tables in the database. However, this process is slow because of the overhead costs
15 incurred by the database management system (particularly relational database management systems). For example, the Oracle Analyze command utilizes the SQL call UPDATE_STATISTICS_FOR TABLE, which invokes all the overhead associated with the Oracle DBMS.
20 Therefore, conventional DBMS analysis are restricted to a few items of information and are relatively slow.

 The Fast Analysis process is significantly faster than the Oracle analyze command. The speed enhancements are achieved by using direct asynchronous
25 I/O, performing block read aheads when reading the database data files and by parallelizing the processing and analysis of data contained in the files.

 By using direct asynchronous I/O to read database files, Fast Analysis is able to directly access
30 specific blocks of a file, thereby bypassing the overhead of typical DBMS routines (SQL shared memory techniques, and other parameters required to access the database tables, for example).

- 24 -

The additional overhead required of typical DBMS calls occurs because SQL and other shared retrieved techniques form a generalized language implemented to provide flexibility in retrieval of database
5 information. In contrast, the Fast Analysis is more singular in purpose and directly retrieves the information needed.

Fast Analysis operates on the following assumptions:

- 10 1. The Fast Analysis process has DBA privileges or read permissions to necessary system files (system dictionary, for example).
2. The Fast Analysis process has read permissions to the data files of the DBMS and
15 particularly any specific tables to be analyzed.

Considering the DBMS Block Structure as discussed above (blocks or pages, extents, and extent header or equivalent structure), the Fast Analysis process is
20 described. Referring to Fig. 4, a user (database operator, DBA, or other mechanism inquiring as to the condition of a database table) first identifies the table of which the condition is to be determined. The table is normally identified by providing the table
25 name (Step 100), and includes determining the location of the extent header segment block for the given table (System Dictionary, for example). The location of the extent segment header block can normally be found in another table that specifies the location of the header
30 block.

At step 105, the extent segment header block is read and an extent list is created. The extent list includes identifying information for each set of blocks

- 25 -

allocated for the table. For example, in one embodiment of the present invention, the extent list includes a file number, block number, and length for each extent in the list. Other types of identifying information such as pages, location or size of the data to be processed, depending on the structure of the data base files, is also appropriate for construction of the extent list. In a more basic description, the extent list need only identify the location(s) of all space allocated for the identified table.

At step 110, the Fast Analysis process determines a number of threads to spawn. The number of threads to spawn can be based upon the number of blocks, or size of the table to be analyzed. This includes reading the extent list to determine how many blocks need to be processed. Physical limitations on the number of useful threads (for example, the number of separate processors available for spawned processes) may be considered. In addition, a specific physical condition or an optimal number of blocks that may be processed in a single thread (100 megabytes of data or 10 blocks, for example) may limit or guide determination of the number of threads to spawn.

Included in step 110 is the creation of one or more sub-lists from the list created at step 105. Each sub-list is preferably created using either an equal number of blocks or proportionately as close to equal as possible to maximize efficiency of the separate parallel processes (threads) to which the lists will be assigned.

If necessary, an item may also be split from the original list, and unequal lists of blocks may be utilized. For example, if 1001 blocks are to be

- 26 -

processed, sublists containing 500/501, 300/300/401, and 50/800/151 may be spawned, the latter combinations only having less efficiency assuming equal processing power in each thread.

5 Each sublist contains the following information for each item in the list: File number, starting block, and number of blocks to process from the starting block number. Again, similar pertinent information may be described in terms of page count, bytes, or other units
10 of data storage and not depart from the scope of the present invention.

 Once the threads are spawned, each thread is assigned a number of blocks to process (step 115). Preferably, the assigned blocks are contained in one of
15 the sublists created in step 110. Alternately, the sublists may be created upon (or contemporaneously with) assignment. Other variations regarding the order or responsibility of each step may be made, again, without departing from the scope or intent of the
20 present invention.

 At step 120, the assigned blocks are read into memory or other storage location using direct asynchronous I/O. As discussed above, the use of direct asynchronous I/O provides direct and immediate
25 access to the data stored in the database table.

 At step 125, the blocks read are examined to determine the contents of each block. In this manner, block and row statistics are gathered from the blocks read. Once all threads have finished processing their
30 respective assigned set of blocks, summation and averaging of information gathered from each thread is performed (Step 130).

- 27 -

At Step 135, the gathered information is either displayed or utilized to update a condition table or provided to another mechanism for using the information gathered. The process may be repeated for any number
5 of tables that may be contained in the data base.

Thus, Fast Analysis provides comprehensive information regarding the condition of database tables and increases the speed at which database information is retrieved and processed. Table 2 illustrates a DDL
10 definition of a DAO repository that may be utilized to store Fast Analysis information extracted from the database tables.

- 28 -

Table 2 - DDL Definition of DAO Repository

```

/
5      CREATE TABLE DAO.COMMON_DICT
      OBJECT_ID      NUMBER
      CONSTRAINT     TABLE_DICT_ID
      PRIMARY KEY
      USING INDEX     TABLESPACE DA02,
10     COLLECT_ID     NUMBER
      CONSTRAINT     TABLE_DICT_COLLECT_ID
      REFERENCES      DAO.COLLECT,
      PCT FREE        NUMBER,
      PCT USED        NUMBER,
      CACHE           NUMBER(1),
15     TABLE LOCK    VARCHAR2(8),
      NESTED          NUMBER(1),
      NUM ROWS        NUMBER,
      BLOCKS          NUMBER,
      EMPTY BLOCKS    NUMBER,
20     AVG SPACE       NUMBER,
      CHAIN CNT       NUMBER,
      AVG ROW LEN     NUMBER,
      AVG SPACE FREELIST BLOCKS NUMBER,
      NUM FREELIST BLOCKS NUMBER,
25     SAMPLE_SIZE     NUMBER,
      IOT NAME        VARCHAR2(30),
      IOT TYPE        NUMBER(1),
      TABLE TYPE OWNER VARCHAR2(30),
      TABLE TYPE     VARCHAR2(30),
30     PACKED          NUMBER(1),
      MIN EXTENT SIZE  NUMBER,
      MAX EXTENT SIZE  NUMBER,
      AVG EXTENT SIZE  NUMBER,
      STD EXTENT_SIZE  NUMBER
35  /
      CREATE TABLE DAO.TABLE_FAST_ANAL
      TABLE_ID      NUMBER
      CONSTRAINT     TABLE_FAST_ANAL_ID
      PRIMARY KEY
40     USING INDEX TABLESPACE DA02,
      TRUE CHAIN CNT  NUMBER,
      TRUE MIGRATE CNT NUMBER,
      FREE LIST_LENGTH NUMBER,
      EMB CNT         NUMBER,
45     SEGMENT _HEADER _HIGH_WATER NUMBER,
      BLOCK _HEADER SIZE NUMBER,
      BLOCK DATA BYTES NUMBER,
      MIN _NORMAL ROWS BYTES NUMBER,
      MAX _NORMAL ROWS BYTES NUMBER,
50     AVG _NORMAL ROWS BYTES NUMBER,
      STI) NORMAL ROWS BYTES NUMBER,
      MIN DELETE ROWS BYTES NUMBER,
      MAX DELETE ROWS BYTES NUMBER,
      AVG DELETE ROWS BYTES NUMBER,
55     STD DELETE ROWS BYTES NUMBER,
      MIN MIGRATED ROWS BYTES NUMBER,
      MAX MIGRATED ROWS BYTES NUMBER,
      AVG MIGRATED ROWS BYTES NUMBER,
      STD _MIGRATED ROWS BYTES NUMBER,
60     MIN CHAINED ROW PIECE BYTES NUMBER,
      AVG CHAINED ROW PIECE BYTES NUMBER,
      STD CHAINED ROW PIECE BYTES NUMBER,
      MIN NEW ROW PIECE BYTES NUMBER,
      MAX NEW ROW PIECE BYTES NUMBER,
65     AVG_NEW ROW_PIECE_BYTES NUMBER,
      REFERENCES      DAO.COLLECT
/

```

- 29 -

Table 2 is not intended to be a working copy, but merely an example repository structure. Fast Analysis may be conveniently implemented using a conventional general purpose or a specialized digital computer or microprocessor programmed according to the teachings of the present specification, as will be apparent to those skilled in the computer art.

Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent to those skilled in the software art. The invention may also be implemented by the preparation of application specific integrated circuits or by interconnecting an appropriate network of conventional component circuits, as will be readily apparent to those skilled in the art.

The present invention includes a computer program product which is a storage medium (media) having instructions stored thereon/in which can be used to program a computer to perform any of the processes of the present invention. The storage medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions.

Stored on any one of the computer readable medium (media), the present invention includes software for controlling both the hardware of a general purpose or specialized computer and for enabling the computer to interact with a human user or other mechanism utilizing the product of the present invention. Such software may include, but is not limited to, device drivers,

- 30 -

operating systems, and user applications. Ultimately, such computer readable media further includes software for determining the condition of a database or for specific database tables, as described above.

5 Included in the programming (software) of the general purpose or specialized computer are software modules for implementing the teachings of the present invention, including, but not limited to, identification and retrieval of database table
10 structures, reading database files and tables, performing analysis of data retrieved from database tables and other related information, and the display, storage, or communication of results as determined by the present invention.

Obviously, numerous modifications and variations of the present invention are possible in light of the above teachings. It is therefore to be understood that within the scope of the appended claims, the invention may be practiced otherwise than as specifically described herein.

- 31 -

CLAIMS

What is claimed is:

1 1. A method of determining a condition of a
2 database, comprising the steps of:
3 identifying a database table;
4 determining elements corresponding to the
5 condition of the database table based on data contained
6 in the database table; and
7 providing said elements to one of a display for
8 viewing by a user, a storage location, and a mechanism
9 for using said elements;
10 wherein said elements include at least one of
11 number of simple rows, number of migrated rows, number
12 of deleted rows, number of blocks initialized but
13 empty, average free space in a block unused, average
14 free released free space, average row length of
15 migrated rows, average row length of chained rows,
16 number of blocks to be read for normal, chained, and
17 migrated rows, measure of density of blocks used,
18 number of blocks exceeding a percentage used threshold,
19 number of blocks less a percentage free threshold,
20 number of chained rows spanning data files, and min max
21 average of chained rows spanning blocks.

1 2. The method according to Claim 1, further
2 comprising the step of repeating said steps of
3 identifying, determining, and providing for each of
4 plural database tables.

1 3. The method according to Claim 1, wherein said
2 step of identifying comprises the steps of:

- 32 -

3 finding a location of an extent segment header
4 block for the database table;
5 reading the extent segment header block; and
6 creating an extent list identifying each extent of
7 the database table based on the segment header block.

1 4. The method according to Claim 3, wherein said
2 step of determining comprises the steps of:
3 gathering data related to said elements from each
4 of data blocks identified in each extent contained in
5 the extent list; and
6 determining said elements via an analysis of the
7 data gathered.

1 5. The method according to Claim 3, wherein said
2 step of creating includes the step of identifying a
3 file number, block number, and length for each extent
4 of the database table.

1 6. The method according to Claim 3, wherein said
2 step of determining comprises the steps of:
3 dividing the extent list into at least two
4 sublists each containing a number of data blocks;
5 spawning a thread for working each of said
6 sublists;
7 gathering information related to said elements by
8 each thread and respective sublist; and
9 performing analysis including summation and
10 averaging of the information gathered by each thread
11 and respective sublist.

1 7. The method according to Claim 6, wherein said
2 step of dividing comprises:

- 33 -

3 dividing the extent list into at least two
4 sublists each containing an equal number of data
5 blocks.

1 8. The method according to Claim 6, wherein said
2 step of gathering includes the steps of:
3 reading each of data blocks identified in each
4 sublist using direct asynchronous I/O; and
5 examining the contents of each data block to
6 determine information related to each of said elements.

1 9. The method according to Claim 8, further
2 comprising the step of repeating said steps of
3 identifying and determining for each of plural database
4 tables.

1 10. A method of analysis of a database,
2 comprising the steps of:
3 identifying a table in said database;
4 reading each of blocks associated with said table
5 using direct asynchronous I/O; and
6 performing said analysis based on information
7 contained in said blocks.

1 11. The method according to Claim 10, further
2 comprising the step of:
3 repeating said steps of identifying, reading, and
4 performing for each of plural tables of said database.

1 12. The method according to Claim 10, wherein
2 said step of reading includes the steps of:
3 creating plural block lists, each list identifying
4 a set of blocks associated with said table;

- 34 -

5 spawning plural threads; and
6 performing the following steps for each respective
7 thread,
8 assigning one of said plural block lists to
9 the respective thread,
10 reading each block identified by the assigned
11 block list using direct asynchronous I/O; and
12 gathering information for said analysis
13 utilizing data contained in each block read.

1 13. The method according to Claim 10, wherein
2 said step of identifying comprises the steps of:
3 determining the location of an extent segment
4 header block of a database table;
5 reading the extent segment header block; and
6 creating an extent list identifying each extent
7 contained in the extent segment header block by file
8 number, block number, and length.

1 14. The method according to Claim 13, wherein:
2 said step of reading each of blocks associated
3 with said table includes the steps of,
4 dividing the extent list into at least two
5 sublists each containing a number of blocks,
6 spawning a thread for working each of said
7 sublists, and
8 gathering information related to said
9 elements by each thread; and
10 said step of performing includes the step of
11 summing and averaging the information gathered by each
12 thread to perform said analysis.

- 35 -

1 15. The method according to Claim 12, further
2 comprising the step of assigning each thread to a
3 separate processor.

1 16. The method according to Claim 14, further
2 comprising the step of assigning each thread to a
3 separate processor.

1 17. The method according to Claim 14, wherein
2 said step of dividing comprises:
3 dividing the extent list into at least two
4 sublists each containing an equal number of blocks.

1 18. A computer readable medium having computer
2 instructions stored thereon that, when loaded into a
3 computer, cause the computer to perform the steps of:
4 identifying a database table;
5 determining elements corresponding to the
6 condition of the database table based on data contained
7 in the database table; and
8 providing said elements to one of a display for
9 viewing by a user, a storage location, and a mechanism
10 for using said elements;
11 wherein said elements include at least one of
12 number of simple rows, number of migrated rows, number
13 of deleted rows, number of blocks initialized but
14 empty, average free space in a block unused, average
15 free released free space, average row length of
16 migrated rows, average row length of chained rows,
17 number of blocks to be read for normal, chained, and
18 migrated rows, measure of density of blocks used,
19 number of blocks exceeding a percentage used threshold,
20 number of blocks less a percentage free threshold,

- 36 -

21 number of chained rows spanning data files, and min max
22 average of chained rows spanning blocks.

1 19. The computer readable medium according to
2 Claim 18, wherein said computer instructions further
3 cause the computer to perform the step of:
4 repeating said steps of identifying, determining,
5 and providing for each of plural database tables.

1 20. The computer readable medium according to
2 Claim 18, wherein said step of identifying comprises
3 the steps of:
4 finding a location of an extent segment header
5 block for the database table;
6 reading the extent segment header block; and
7 creating an extent list identifying each extent of
8 the database table based on the segment header block.

1 21. The computer readable medium according to
2 Claim 20, wherein said step of determining comprises
3 the steps of:
4 gathering data related to said elements from each
5 of data blocks identified in each extent contained in
6 the extent list; and
7 determining said elements via an analysis of the
8 data gathered.

1 22. The computer readable medium according to
2 Claim 20, wherein said step of creating includes the
3 step of identifying a file number, block number, and
4 length for each extent of the database table.

- 37 -

1 23. The computer readable medium according to
2 Claim 20, wherein said step of determining comprises
3 the steps of:
4

dividing the extent list into at least two

- 37 -

1 23. The computer readable medium according to
2 Claim 20, wherein said step of determining comprises
3 the steps of:

4 dividing the extent list into at least two
5 sublists each containing a number of data blocks;

6 spawning a thread for working each of said
7 sublists;

8 gathering information related to said elements by
9 each thread and respective sublist; and

10 performing analysis including summation and
11 averaging of the information gathered by each thread
12 and respective sublist.

1 24. The computer readable medium according to
2 Claim 23, wherein said step of dividing comprises:

3 dividing the extent list into at least two
4 sublists each containing an equal number of data
5 blocks.

1 25. The computer readable medium according to
2 Claim 23, wherein said step of gathering includes the
3 steps of:

4 reading each of data blocks identified in each
5 sublist using direct asynchronous I/O; and

6 examining the contents of each data block to
7 determine information related to each of said elements.

1 26. The computer readable medium according to
2 Claim 25, wherein said computer instructions further
3 cause the computer to perform the step of:

4 repeating said steps of identifying and
5 determining for each of plural database tables.

- 38 -

1 27. A computer readable medium having computer
2 instructions stored thereon that, when loaded into a
3 computer, cause the computer to perform the steps of:
4 identifying a table in said database;
5 reading each of blocks associated with said table
6 using direct asynchronous I/O; and
7 performing said analysis based on information
8 contained in said blocks.

1 28. The computer readable medium according to
2 Claim 27, wherein said computer instructions further
3 cause the computer to perform the step of:
4 repeating said steps of identifying, reading, and
5 performing for each of plural tables of said database.

1 29. The computer readable medium according to
2 Claim 27, wherein said step of reading includes the
3 steps of:
4 creating plural block lists, each list identifying
5 a set of blocks associated with said table;
6 spawning plural threads; and
7 performing the following steps for each respective
8 thread,
9 assigning one of said plural block lists to
10 the respective thread,
11 reading each block identified by the assigned
12 block list using direct asynchronous I/O; and
13 gathering information for said analysis
14 utilizing data contained in each block read.

1 30. The computer readable medium according to
2 Claim 27, wherein said step of identifying comprises
3 the steps of:

- 39 -

4 determining the location of an extent segment
5 header block of a database table;
6 reading the extent segment header block; and
7 creating an extent list identifying each extent
8 contained in the extent segment header block by file
9 number, block number, and length.

1 31. The computer readable medium according to
2 Claim 30, wherein:
3 said step of reading each of blocks associated
4 with said table includes the steps of,
5 dividing the extent list into at least two
6 sublists each containing a number of blocks,
7 spawning a thread for working each of said
8 sublists, and
9 gathering information related to said
10 elements by each thread; and
11 said step of performing includes the step of
12 summing and averaging the information gathered by each
13 thread to perform said analysis.

1 32. The computer readable medium according to
2 Claim 29, wherein said computer instructions further
3 cause the computer to perform the step of assigning
4 each thread to a separate processor.

1 33. The computer readable medium according to
2 Claim 31 wherein said computer instructions further
3 cause the computer to perform the step of assigning
4 each thread to a separate processor.

1 34. The computer readable medium according to
2 Claim 31, wherein

- 40 -

3 said step of dividing comprises:
4 dividing the extent list into at least two
5 sublists each containing an equal number of blocks.

1 35. An apparatus for determining a condition of
2 a database, comprising:
3 means for identifying a database table;
4 means for determining elements corresponding to
5 the condition of the database table based on data
6 contained in the database table; and
7 means for providing said elements to one of a
8 display for viewing by a user, a storage location, and
9 a mechanism for using said elements;
10 wherein said elements include at least one of
11 number of simple rows, number of migrated rows, number
12 of deleted rows, number of blocks initialized but
13 empty, average free space in a block unused, average
14 free released free space, average row length of
15 migrated rows, average row length of chained rows,
16 number of blocks to be read for normal, chained, and
17 migrated rows, measure of density of blocks used,
18 number of blocks exceeding a percentage used threshold,
19 number of blocks less a percentage free threshold,
20 number of chained rows spanning data files, and min max
21 average of chained rows spanning blocks.

1 36. The apparatus according to Claim 35, wherein
2 said means for identifying comprises:
3 means for finding a location of an extent segment
4 header block for the database table;
5 means for reading the extent segment header block;
6 and

- 41 -

7 means for creating an extent list identifying each
8 extent of the database table based on the segment
9 header block.

1 37. The apparatus according to Claim 35, wherein
2 said means for determining comprises:
3 means for gathering data related to said elements
4 from each of data blocks identified in each extent
5 contained in the extent list; and
6 means for determining said elements via an
7 analysis of the data gathered.

1 38. The apparatus according to Claim 36, wherein
2 said means for determining comprises:
3 means for dividing the extent list into at least
4 two sublists each containing a number of data blocks;
5 means for spawning a thread for working each of
6 said sublists;
7 means for gathering information related to said
8 elements by each thread and respective sublist; and
9 means for performing analysis including summation
10 and averaging of the information gathered by each
11 thread and respective sublist.

1 39. The apparatus according to Claim 38, wherein
2 said means for gathering includes:
3 means for reading each of data blocks identified
4 in each sublist using direct asynchronous I/O; and
5 means for examining the contents of each data
6 block to determine information related to each of said
7 elements.

- 42 -

1 40. An apparatus for analysis of a database,
2 comprising:
3 means for identifying a table in said database;
4 means for reading each of blocks associated with
5 said table using direct asynchronous I/O; and
6 means for performing said analysis based on
7 information contained in said blocks.

1 41. The apparatus according to Claim 40, wherein
2 said means for reading includes:
3 means for creating plural block lists, each list
4 identifying a set of blocks associated with said table;
5 means for spawning plural threads; and
6 means for performing the following steps for each
7 respective thread,
8 assigning one of said plural block lists to
9 the respective thread,
10 reading each block identified by the assigned
11 block list using direct asynchronous I/O; and
12 gathering information for said analysis
13 utilizing data contained in each block read.

1 42. The apparatus according to Claim 40, wherein
2 said means for identifying comprises:
3 means for determining the location of an extent
4 segment header block of a database table;
5 means for reading the extent segment header block;
6 and
7 means for creating an extent list identifying each
8 extent contained in the extent segment header block by
9 file number, block number, and length.

1 43. The apparatus according to Claim 40, wherein:

- 43 -

2 said means for reading each of blocks associated
3 with said table includes,
4 means for dividing the extent list into at
5 least two sublists each containing a number of
6 blocks,
7 means for spawning a thread for working each
8 of said sublists, and
9 means for gathering information related to
10 said elements by each thread; and
11 said means for performing includes means for
12 summing and averaging the information gathered by each
13 thread to perform said analysis.

1 44. The apparatus according to Claim 43, wherein
2 said means for
3 dividing comprises:
4 means for dividing the extent list into at least
5 two sublists each containing an equal number of blocks.

1 45. An apparatus for determining a condition as
2 of a database, comprising:
3 an identifying device configured to identify a
4 database table;
5 a determining device configured to determine
6 elements corresponding to the condition of the database
7 table based on data contained in the database table;
8 and
9 a mechanism for using said elements;
10 wherein said elements include at least one of
11 number of simple rows, number of migrated rows, number
12 of deleted rows, number of blocks initialized but
13 empty, average free space in a block unused, average
14 free released free space, average row length of

- 44 -

15 migrated rows, average row length of chained rows,
16 number of blocks to be read for normal, chained, and
17 migrated rows, measure of density of blocks used,
18 number of blocks exceeding a percentage used threshold,
19 number of blocks less a percentage free threshold,
20 number of chained rows spanning data files, and min max
21 average of chained rows spanning blocks.

1 46. The apparatus according to Claim 45, wherein
2 said identifying device comprises:

3 a locator configured to locate an extent segment
4 header block for the database table;

5 a read mechanism configured to read the extent
6 segment header block; and

7 a list mechanism configured to create an extent
8 list identifying each extent of the database table
9 based on the segment header block.

1 47. The method according to Claim 45, wherein
2 said determining device comprises:

3 a divider configured to divide the extent list
4 into at least two sublists each containing a number of
5 data blocks;

6 a process spawner configured to a thread for
7 working each of said sublists;

8 an information gatherer configured to information
9 related to said elements by each thread and respective
10 sublist; and

11 an analyzer configured to perform an analysis
12 including summation and averaging of the information
13 gathered by each thread and respective sublist.

- 45 -

14 dividing the extent list into at least two
15 sublists each containing an equal number of data
16 blocks.

1 48. The method according to Claim 47, wherein
2 said information gatherer comprises:

3 a reader configured to read each of data blocks
4 identified in each sublist using direct asynchronous
5 I/O; and

6 an examiner configured to examine the contents of
7 each data block to determine information related to
8 each of said elements.

1 49. An apparatus for analysis of a database,
2 comprising:

3 an identifying device configured to identify a
4 table in said database;

5 a reader configured to read each of blocks
6 associated with said table using direct asynchronous
7 I/O; and

8 an analyzer configured to perform said analysis
9 based on information contained in said blocks.

1 50. The apparatus according to Claim 49, wherein
2 said reader comprises:

3 a list device configured to create plural block
4 lists, each list identifying a set of blocks associated
5 with said table; and

6 a process spawner configured to spawn threads for
7 processing and assign one of said block lists to each
8 thread spawned; and

9 wherein the following steps are performed by each
10 respective thread,

- 46 -

11 reading each block identified by the assigned
12 block lists using direct asynchronous I/O; and
13 gathering information for said analysis
14 utilizing data contained in each block read.

1 51. The apparatus according to Claim 49, wherein
2 said identifying device comprises:
3 a locator configured to determine a location of an
4 extent segment header block of the database table and
5 read the extent segment header block; and
6 a list device configured to create an extent list
7 identifying each extent contained in the extent segment
8 header block.

1 52. The method according to Claim 49, wherein:
2 said reader comprises,
3 a divider configured to divide the extent
4 list into at least two sublists each containing a
5 number of blocks,
6 a process spawner configured to spawn a
7 thread for working each of said sublists, and
8 an information gatherer configured to gather
9 information related to said elements by each
10 thread; and
11 said analyzer comprises a summing device
12 configured to sum and average the information gathered
13 by each thread to perform said analysis.

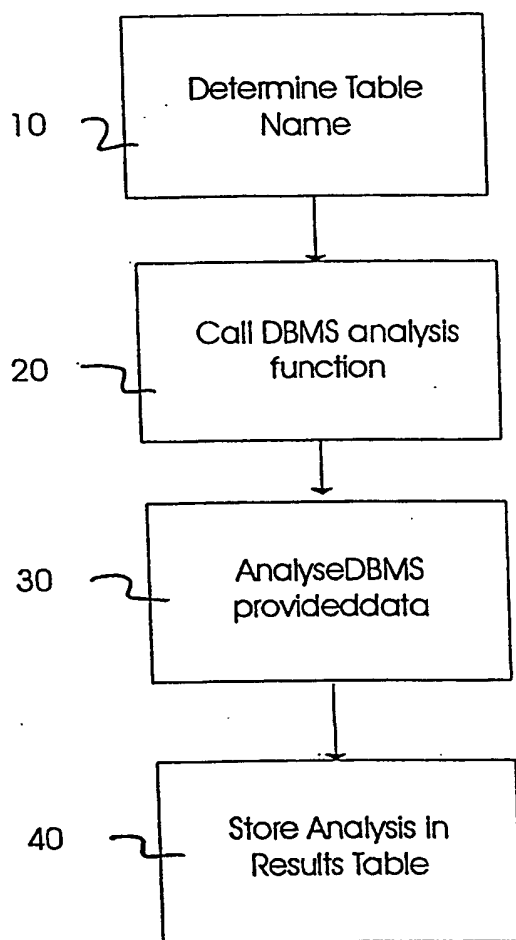


Fig. 1

Prior Art

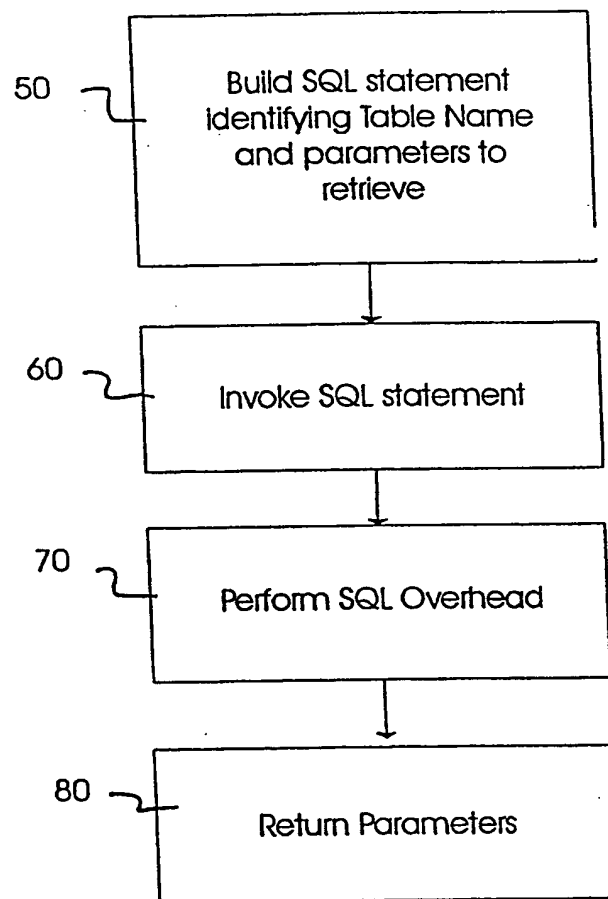


Fig. 2

Prior Art

3/4

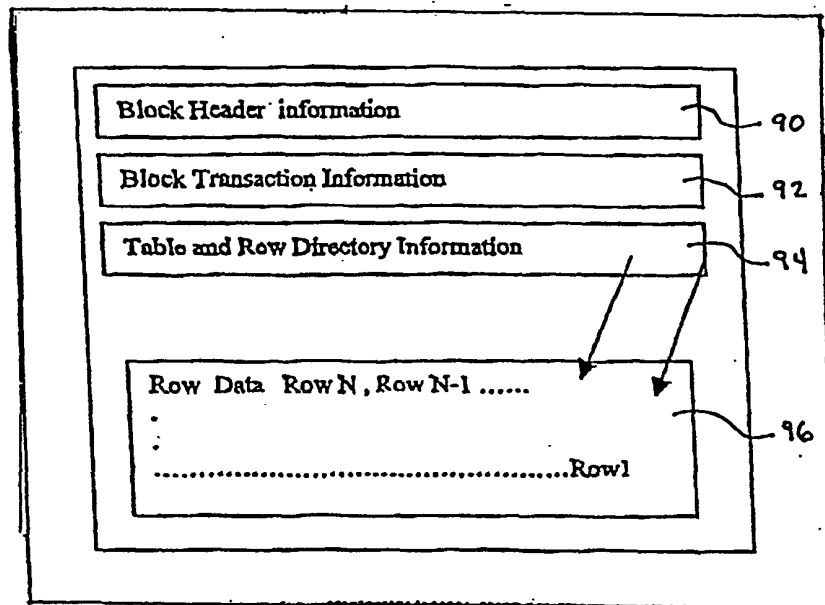


Fig. 3

4/4

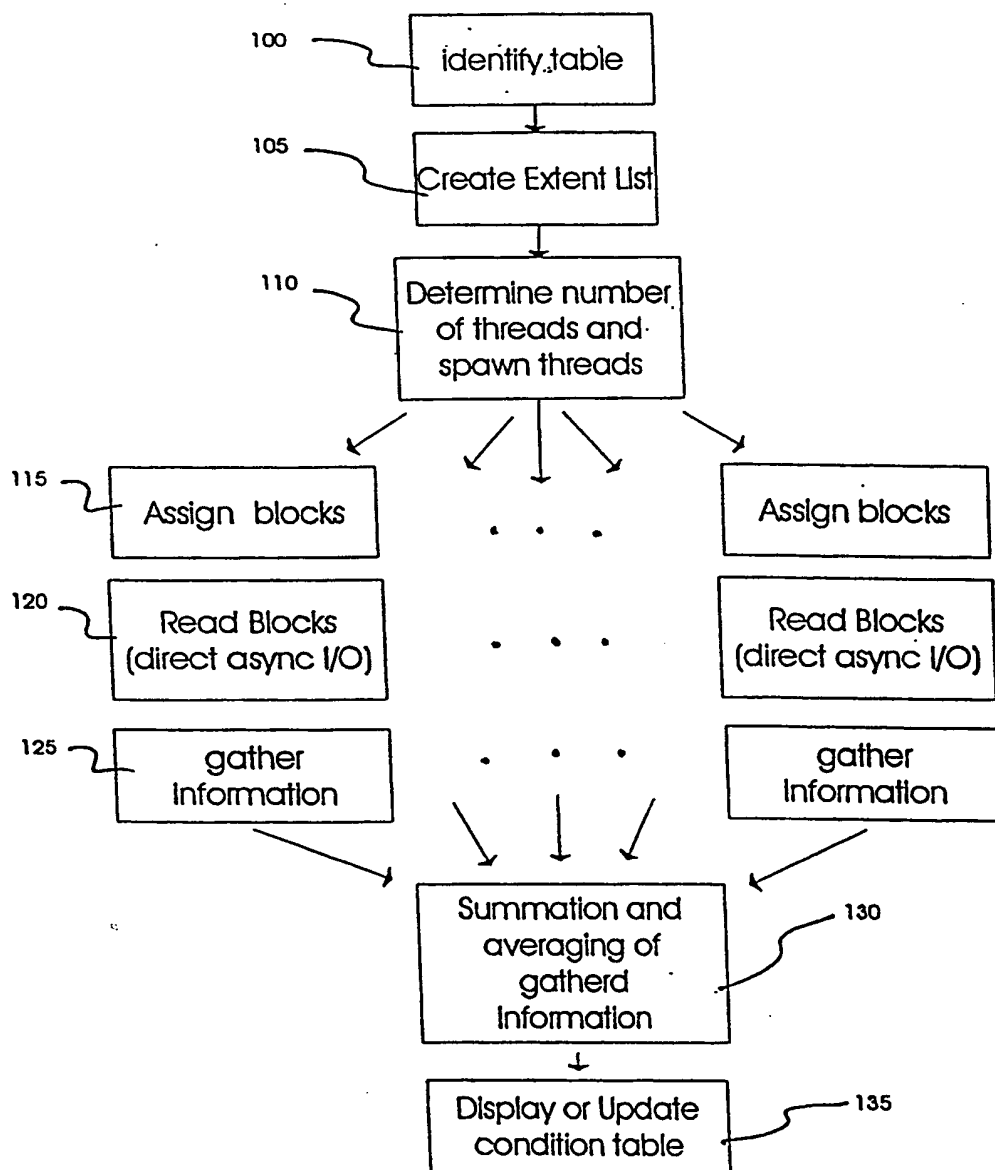


Fig. 4